

Spot On - Understanding Halftones

DTP Reviews
DTP Articles
DTP Tutorials

[[Understanding Halftones](#)]

[[XML Revolution](#)]

[[Windows 2000](#)]

[[PDF Workflows](#)]

[[Design Developers](#)]

[[Text Typography](#)]

[[Spot Colour](#)]

[[Database Publishing](#)]

[[Layout Compositing](#)]

[[Design Automation](#)]

[[Stock Photography](#)]

[[Quark v Adobe](#)]

[[Asymmetric Grids](#)]

[[Commercial Print](#)]

[[Preparing Print](#)]

[[Understanding EPS](#)]

[[Layout Grids](#)]

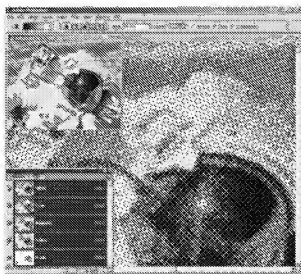
[[Beyond CYMK](#)]

[[Acrobat v Immedia](#)]

[[Postscript 3](#)]

Halftoning - the secret print ingredient

Tom Arah looks at the crucial role of halftone screening in successful print.

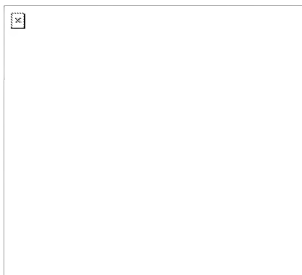


Last month I probably upset a lot of Web designers by saying that much of their work boils down to producing buttons. This month I'm going to offend all print designers by showing how all their work ultimately comes down to handling dots.

At first sight this might seem obvious. After all everyone knows that print is output as a grid of dots, that's where the dots-per-inch (dpi) of your laser, inkjet or imagesetter come in. The higher your device's output resolution, the more dots can be switched on or off, black or white, resulting in higher quality serifs, crisper lines and so on. The input side for scanning black or white line-art is just as straightforward. The more dpi that your scanner is capable of resolving, the more detail can be captured (though the limitations of the eye mean that there's no point scanning at higher than 1000 dpi size-for-size).

For bilinear black or white work, such as text-only work or outputting architectural plans, there's a relatively straightforward mapping from monotone input dots to monotone output dots. However the world isn't monotone, it's made up of continuous tones (and colour which we'll come onto later) and it's here that this simple mapping breaks down. For input this isn't really an issue as each scanned pixel isn't limited to representing a bilinear on-or-off dot but can just as easily represent 256 greyscales or 16 million RGB values

(which is why rather than dpi, scanner resolutions should really be measured in spi or samples per inch).



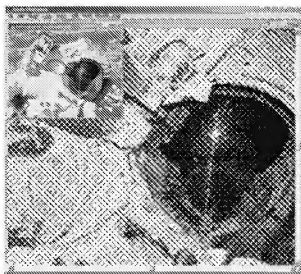
There's a fundamental difference between bilinear black or white work and continuous tone grayscale and colour.

Output though is another matter. The toner in your laser or the ink on the offset press is limited to just one colour - usually black - so how can it produce continuous tones, the full spectrum of tints of gray? This has been a central issue for the printing industry since it was invented and the solution is to rely on an optical illusion. Engravers realized that by using closely drawn cross-hatching, the eye isn't aware of the individual lines but instead averages the effect to produce the impression of gray. The closer and the heavier the hatching, the darker the perceived percentage tint.

For centuries this manual artistic system sufficed, but with the advent of photography the need for an automatic photo-mechanical solution became essential. Around 1870 the solution emerged - halftoning. This halftoning process works by breaking down the continuous tone image into a grid of regularly spaced cells within which are centred monotone but variably sized "halftone spots". The larger the spots, the darker the perceived gray.

The principle is clear but how can the effect be created? The solution developed was "screening" which involves re-photographing the continuous tone image through a variable density mesh onto a high contrast, black or white film. The amount of light reflected for each cell of the screen determines the density of the mesh that can be penetrated so that highlights produce larger spots and shadows produce smaller spots on the resulting film negative. This negative is then

reversed on the press to produce the desired effect of larger halftone spots in shadow areas and smaller spots in highlights.



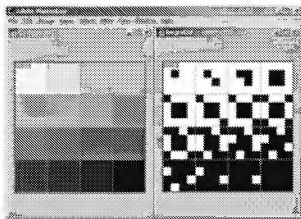
Halftoning breaks down the continuous tone into regularly spaced but variably sized spots.

Only fifteen years ago these process cameras were the one piece of equipment that a graphic design studio had to possess. Nowadays of course the camera has been replaced by the computer, but the same halftone screening process underpins today's all-digital approach by converting a pixel's gray or colour levels into varying sizes of halftone spot. And it is these halftone spots rather than the underlying device dots that are the secret behind just about every image, tint or colour that you see in print!

As such it's clearly a good idea to understand exactly what is involved in the digital halftoning process and which parameters are important to its success. Much the most important factor is the "line frequency" or "screen ruling" of the halftone screen. It is this setting, measured in lines per inch (lpi), that determines the size and spacing of the cells in the halftone grid. It's desirable for this frequency to be high as this means that the halftone spots are small and close together so less discernible to the eye. While a newspaper might use a line ruling of 85 lpi, for example, higher quality print would use 133 or 150 lpi.

In which case why not use a high screen frequency of 150 lpi for all print? This brings us back to those dots. Always remember that a halftone spot is not a device dot - but it is composed of them. With a 600 dpi printer and a 150 lpi screen frequency this means that each halftone cell is made up of a square of 4 x 4 device dots. This means that each halftone spot can only represent 16 different gray levels ranging from

all device dots in the cell being off/white to all device dots being on/black (OK there are actually 17 possible levels if you're being pedantic). The end result for our image would be severe and unacceptable posterisation.



The number of pixels in each halftone cell determine the number of possible gray scales.

We clearly need to be able to produce more than 17 gray levels and the commonly accepted industry standard is 256, as this is more than the average press or eye can discriminate and is also the maximum number that Postscript Levels 1 and 2 can produce (Level 3 can produce 4096 levels which makes a difference for gradients). To create these optimal 256 levels we therefore need a grid of 16×16 device dots. For imagesetters that's not too much of a problem as you can simply multiply the desired ruling by sixteen to determine the necessary resolution so that for quality work set at 150 lpi you need 2400 dpi output while for 85 lpi newspaper work you can get away with 1400 dpi. For a 600 dpi laser printer though there has to be a trade-off with a typical compromise being a screen frequency of 60 lpi resulting in a 10×10 halftone cell capable of producing 101 gray levels.

This fundamental distinction between halftone spots and device dots has a very beneficial knock-on effect for input as the resolution of your continuous tone scan is determined not by your output device resolution but rather by your output screen frequency. That's why even if your photograph is being output on an imagesetter you don't need to scan it at 1200 or 2400 dpi! In fact as you'll rarely use halftone screenings above 150 lpi you could get away with 150 dpi scans. However this would involve a one-to-one mapping between each pixel and halftone spot which is a bit mean so the usual rule is to scan in your photo at twice the desired output screen ruling. This means that each halftone spot is determined by averaging four (2×2) pixels. More importantly it means that for size-for-size

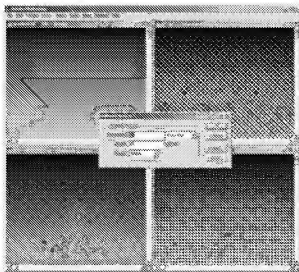
work you never need to scan your photographs at more than 300 dpi!

At this stage, with the principles in place, it's probably a good idea to take a look at halftoning in action and to see some of the issues this throws up. This is easy enough if you start off with a grayscale image in Photoshop (I converted a copy of the CMYK sample astronaut image and also created a black to white gradient of a similar size). If you then convert the grayscale images to monotone bitmaps using the Image > Mode command you are presented with a dialog in which you can choose to halftone the image and a sub dialog in which you can set the ruling, angle and shape. Normally you want to minimize the visibility of the halftone but we want to make it clear so we can set a crude screen ruling of 12 lpi, an angle of 90° and a square dot shape. The result is a solid black and white image but which clearly gives the appearance of the original's grayscales. Zoom in and the halftone cells and spots producing the effect become even more apparent.

We immediately hit a problem however as the solid black shadows in the astronaut's helmet stand out like a sore thumb from the rest of the halftone. This is even more apparent in the gradient image and just as true of the solid white. Unfortunately it's a problem that will be exaggerated on the press where the smallest halftone spots are likely to be "blown out" and lost completely, whereas the larger spots are likely to suffer ink spread resulting in "dot gain" ("spot gain" in a rational world) so that the standout areas of solid white and black will both increase.

The ideal solution is to use a high quality press and paper stock and to fine-tune your image across its tonal range to take into account spot variation. There's one step that is essential for all press work however which is to control your spot range. Essentially this means using your photo editor's Levels or Curves command to darken highlights and to lighten shadows in your scans. The extent of this depends on your paper with the tonal range for images destined for high quality coated stock reduced to between 5% and 95% and for newsprint down to between 12% and 88%! Clearly this cuts the overall detail and contrast in your image but the results on paper will benefit.

That's tackled the problem of solid black and white but there's still plenty to do to improve the quality of the intermediate gray levels. To begin with we want to make the halftone screen as unobtrusive as possible. The eye is particularly sensitive to vertical and horizontal lines so the apparent quality of the image immediately improves if we angle the screen to 45°.



Photoshop lets you explore the effect of different screen frequencies, angles and spot shapes.

Another problem is that the square shape of our halftone spot is very noticeable. Of course at the much higher screen frequencies used for print it wouldn't stand out like this, but it still affects the overall look of the photographic image. It also leads to another noticeable problem, especially apparent in the gradient image, with an "optical jump" where the supposedly smooth spectrum suddenly becomes much darker. This is caused in shadow areas when the halftone spots begin to touch each other. A serious improvement occurs if we change the halftone spot shape to an ellipse as its double axis results in two smaller and less noticeable optical jumps.

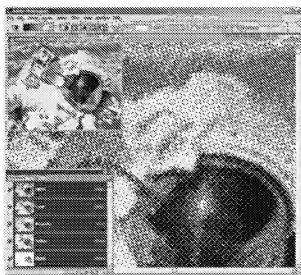
In fact you don't normally need to explicitly set the halftone spot shape at all, except for special effects, as this is better left to Postscript and the imagesetter. This is especially the case in the digital age as there's no reason that the overall spot shape must stay the same whatever its size. Nowadays this shape flexibility can be used to advantage with Postscript defaulting to a "transforming ellipse" where the ellipse starts off as a black spot in a white cell and then switches at 50% to become a white spot in a black cell. In the future, as halftoning algorithms improve, the spot shape could change radically across the tonal range especially when dealing with overlaid halftone screens.

So why would you want to overlay halftone screens? The answer of course is colour. So far I've only talked about simulating grayscales but the real beauty of halftones is that they can easily be adapted to produce colour. The system is analogous to the screen pixel where 256 levels of red, green and blue lead to 16 million possible colours ($256 \times 256 \times 256$).

With print the primaries are cyan, magenta and yellow. By overlaying angled halftone screens printed in these inks the offset halftone spots create a tiny rosette pattern that the eye again averages to produce the impression of a colour. This is the secret behind colour separation and all colour print.

Of course it's not quite as simple as this. To begin with there's not a one-to-one mapping between the RGB and CMY colour spaces (though that's another story). More importantly although 100% cyan, magenta and yellow should produce black in fact it's more of a very dark brown, subject to over-inking and also limited to the screen ruling rather than the output resolution which is very undesirable for text readability. The solution of course is to add an extra dedicated black separation and it is this four-plate, CMYK system that the publishing industry is built on.

Again it's a good idea to see the process in action and again Photoshop comes in handy. By opening the sample astronaut image in its native CMYK we can then apply the Pixelate>Colour Halftone filter with a 4 pixel radius. The overall result is a recognizable colour halftone version of the original image but looking in the highlight areas it is also clear how the effect is created with clearly discernible discrete CMY and K halftone spots offset to each other to produce the typical rosette pattern of process print. Look at any commercially printed colour photograph or tint under a magnifying glass and you'll see exactly the same effect. Inspecting the separate CMYK channels in turn we can see exactly how the effect is produced - effectively the bitmap filter we looked at earlier has been applied to each channel in turn with the same ruling and shape but with a different angle.



Photoshop's colour halftone filter shows how CMYK

halftoning produces the impression of continuous tone colour.

But which angle? With just three colours it's easy to angle each plate at 30 degrees to each other (grid geometry means we've only got 90° to play with not 360°), but with four plates that's not an option. The solution that the industry has standardized on, is to keep the darkest plate, black, on the least noticeable 45° angle, to set the lightest colour, yellow, to 0° degrees and to set the cyan and magenta plates 30° away from black on the geometrically equivalent 15° and 75°. In fact there's a little leeway (though absolute accuracy is essential) with Adobe, for example, preferring 71.5651° and 18.4349°!

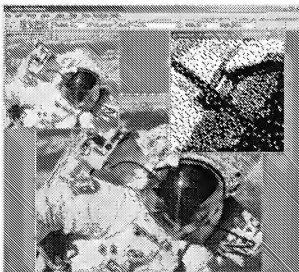
For the majority of commercial print this is the halftoning process employed, but what happens if you need to print extra plates? For most spot colour "bump" plates you can set the angle to 45° as it's rare to combine the spot colour with tints of black. If you do however, for example to produce duotone images, the spot colour must be set to one of the other angles. With extra bump plates of course the situation gets worse. And there's another serious complicating factor. If any screen is not perfectly registered the halftone spots will clash, the rose pattern will disappear and there will be noticeable interference or "moiré".

The potential moiré problem is even worse in the case of "hi-fi" colour where instead of the traditional four CMYK plates there can be up to eight plates. The most common example of hi-fi colour is Pantone's Hexachrome which as its name suggests uses six colours, with additional green and orange plates set on the same angles as magenta and cyan. If any plate is even 1° out of registration your high quality print-run can be ruined.

Clearly we are hitting the limits of the halftone approach so is there an alternative? The answer is stochastic screening which dispenses with the halftone spot completely. Rather than creating the impression of grays and colours by changing the size of the regularly spaced halftone spot, stochastic screening creates the effect by controlling the distribution of irregularly placed but uniformly sized dots. The more concentrated the dots, the darker the perceived colour.

Again you can see the system in action using Photoshop. If you open or create a grayscale image and then convert it to bilinear Bitmap mode one of the options you are presented with is diffusion dither. The end result is a much closer approximation to the original image than we achieved with halftoning. Zoom in and you will see how the effect has been produced by manipulating single pixels rather than by creating

the grid of halftone spots.



Stochastic screening uses the concentration of dots rather than the size of spot to create the impression of grayscales.

For commercial print the benefits of stochastic screening are huge with no distracting angles or halftone shapes, no moiré or optical jumps, fewer problems with misregistration and spot variation, generally higher quality CMYK print with greater tonal range and detail and a much easier upgrade route to the massive quality gain of hi-fi colour. The system isn't perfect, it's much more difficult to proof accurately and there's a clear graininess to images that can be especially distracting in highlights, but there's little question that it's the screening technology of the future.

Eventually then with stochastic screening the difference between halftone spot and device dot might truly disappear. In the meantime though it's crucial to recognize and understand the central role that the under-appreciated halftone spot plays. Halftoning has been the secret behind commercial print for well over a century and will continue to be so for many years to come.

Tom Arah

March 2001

For further information read the excellent Real World Scanning and Halftones by David Blatner, Glenn Fleischman and Steve Roth.

[\[Home\]](#) [\[Writing\]](#) [\[DTP\]](#) [\[Photo\]](#) [\[Draw\]](#)
[\[Web\]](#)



Copyright 1995-2001, Tom Arah, Designer-Info.com. Please get in contact to let me know what you think about a particular piece or the site in general.

If you've found the site useful please remember to bookmark it and to tell others. Also please consider buying your design apps, upgrades, books, CDs and videos using the site's direct links or these general links [Amazon.com](#) or [Amazon.co.uk](#). This costs you absolutely nothing but I gain a small commission which helps cover the time and effort involved in keeping the site up-to-date. Thank you.